

Amendments to the Specification:

Please replace Paragraph [0006] with the following amended paragraph:

Furthermore, with the increasing popularity of the Internet and "always on" infrastructures, many such computer systems ~~many~~ may be deployed in a data center. For example, a modern data center may have hundreds of system racks, with each system rack having four or more MP systems, as described above. Of course, such data centers need air conditioning systems to remove the heat generated by all these computer systems, and the air conditioning systems themselves consume significant power. In addition, lighting, redundant power subsystems, and security systems all contribute to the power consumption of a data center. When all these factors are taken into account, it is not difficult to see that modern data centers can consume megawatts of electricity and have electric bills that reach thousands of dollars per day.

Please replace Paragraph [0012] with the following amended paragraph:

Finally, it is often desirable to provide a distributed application having redundancy. One term used in the art is "N+1" redundancy. Basically, if N components are needed to provide a service, "N+1" components are provided. If one of the N components fails, the service is gracefully shifted to the redundant "+1" component, and the distributed application continues operating normally. However, "N+1" redundancy also increases power consumption because the "+1" component tends to be "hot". In other [word] words, the redundant component remains powered up waiting for a failure in one of the other components. Accordingly, redundancy also increases the power consumption of a distributed application. Of course, redundancy increases revenue for a business that depends on a distributed application because the availability of the application is increased by minimizing down time.

Please replace Paragraph [0014] with the following amended paragraph:

The present invention provides a system and method for intelligent control of power consumption of distributed services and components, such as those used to implement a distributed [applications] application. The present invention is best implemented on a computer system that provides independent computing elements capable of being powered down or entering a power saving mode, thereby allowing individual services or components to be powered down. Note that the granularity with which the power consumption of a distributed application can be varied is provided by the ability of ~~cause~~ individual host processor cards or other computing elements to enter a power saving mode.

Please replace Paragraph [0018] with the following amended paragraph:

Finally, the third algorithm of the present is a minimal power-consuming redundant computing hardware algorithm that provides "N+1" or greater redundancy for the other host processor cards. Basically, one or more host processor cards can be provided as cold spares. If a current failure or impending failure is detected in one of the other cards, the cold spare card enters normal operation mode from power saving mode. Thereafter, the operating system is loaded, and the components of the distributed application that are hosted by the failing card are initialized and begin operating on the cold spare card. At this point, the components executing on the failing card can be gracefully shut down, if possible, and the failing card can be placed into hot swap mode. Once in hot swap mode, the failing card can be replaced with a replacement card. Note that at this point, the replacement can remain in hot swap/power saving mode and serve as the new cold spare. Alternatively, the replacement card can enter normal operation mode, the components can be moved to the replacement card, and the cold spare can be placed into power saving mode and resume its function as a cold spare.

Please replace Paragraph [0031] with the following amended paragraph:

FIG. 1 is a front perspective view illustrating a server system 100 capable of operating with the present invention. FIG. 2 is a rear perspective view illustrating server system 100. Server system 100 includes panels 102, liquid crystal display (LCD) panels 104A and 104B (collectively referred to as LCD panels 104), backplane 106, chassis 108, and dual redundant power supply units 114A and 114B (collectively referred to as power supply units 114). Panels 102 are attached to chassis 108, and provide protection for the internal components of server system 100. Backplane 106 is positioned near the center of server system 100. Backplane 106 is also referred to as midplane 106. LCD panels 104A and 104B are substantially identical, except for their placement on server system 100. LCD panel 104A is positioned on a front side of server system 100, and LCD panel 104B is positioned on a back side of server system 100. Power supply units 114 are positioned at the bottom of server system 100 and extend from a back side of server system 100 to a front side of server system 100. Power supply units 114 each include an associated cooling fan 304 (shown in block form in FIG. 3). Additional cooling fans 304 may also be positioned behind LCD panel 104B. In one configuration, four chassis cooling fans 304 are used in server system 100. In another configuration, six chassis cooling fans 304 are used. Other numbers and placement of cooling fans 304 may be used. Cooling fans 304 may also be configured in a "N+1" redundant cooling system, where "N" represents the total number of necessary fans 304, and "+1" represents the number of redundant fans 304.

Please replace Paragraph [0085] with the following amended paragraph:

In a typical distributed application for an on-line retailer, assume that more customers will be browsing the product catalog and placing orders than checking inventory and tracking shipments. Therefore, only two host processor cards are

needed to host each of the latter two components. Accordingly, host processor cards 300C and 300J are configured to host inventory component 606, with card 300C hosting inventory component 606A and card 300J hosting inventory component [3606B] 606B. Similarly, host processor cards 300D and 300K are configured to host shipment tracking component 608, with card 300D hosting shipment tracking component 608A and card 300K hosting shipment tracking component 608B.

Please replace Paragraph [0091] with the following amended paragraph:

Algorithm 900 exploits the fact that not all components of a distributed application contribute equally to the revenue stream of a business using the distributed application. In accordance with the present invention, if power consumption must be reduced, components having less of a contribution to revenue (or for some other reason, lower priority) should be suspended to save power before components that having a higher contribution to revenue (or for some other reason, higher priority). With reference to distributed application 600 of FIG. 6, to maintain the revenue stream, it is essential that customers have access to product catalog component 602 to select a product to order, and order processing component 604 to place an order for the product. However, it is less important (although certainly still helpful) to the revenue stream for the customer to be able to confirm that the product is in stock or when it will ship using inventory component 606. Furthermore, it is even less important that the customer be able to track shipments using shipment tracking component 608, since a customer will generally not need this function until after an order has been placed and the revenue generated by the order has been secured.

Please replace Paragraph [0097] with the following amended paragraph:

As mentioned above, FIG. 8 illustrates reduced load power saving algorithm 800. In FIG. 8, algorithm 800 is illustrated as a flowchart 800A that

shows [are] how power can be saved when loads are reduced, and flowchart 800B shows how additional capacity can be added in anticipation of increased or peak loads, in accordance with the present invention.

Please replace Paragraph [0104] with the following amended paragraph:

Flowchart 800B starts at "START" block 814, and control passes to block 816. Block 816 anticipates an impending period of increased or peak demand. Note that it is desirable to add additional capacity before it is actually required. By [during] doing so, distributed application 600 can always service transactions quickly and efficiently. As mentioned above, in a typical distributed application, load levels may vary in a predictable manner, so it is possible to detect an anticipated increase in load by detecting that transactions are increasing along a predicable curve.

Please replace Paragraph [0105] with the following amended paragraph:

Next control passes to block 818, which identifies duplicate instances of components of distributed application 600 that will be needed during the period of increased or peak load. In the example above, it was [be] determined that components 602C, 604C, 606B, and 608B were not needed, so these components [where] were suspended and cards 300L, 300M, 300J and 300K were placed in power saving mode. Now assume that these components are again needed.

Please replace Paragraph [0109] with the following amended paragraph:

FIG. 9 illustrates priority-based power consumption reduction algorithm 900. In FIG. 9, algorithm 900 is illustrated as a flowchart 900A that shows how power can be reduced when it is necessary to reduce power consumption by suspending components in priority order, in accordance with the present

invention. Flowchart 900B shows how suspended components can resume operation when power consumption can be increased, in accordance with the present invention.

Please replace Paragraph [0116] with the following amended paragraph:

Next control passes to block 918, which identifies in priority order, such as contribution to the revenue stream, which components should resume operation. In the example above, it was [be] determined that components 608A and 608B could be suspended, so these components were suspended and cards 300D and 300K were placed in power saving mode. Now assume that these components can resume operation because power supplies are sufficient.

Please replace Paragraph [0124] with the following amended paragraph:

If the card 300 in which the actual or impending failure has been detected is still functional, block 1012 attempts to gracefully suspend or shut down all the components identified at block 1006. However, this may not be possible if the affected card 300 is not responding. Control then passes to block 1014, [where] were the card 300 in which the actual or impending failure has been detected is signaled to enter hot swap mode from normal operation mode. Control then passes to "STOP" block 1016. At this point, the card 300 in which the actual or impending failure has been detected can be removed and replaced by a functioning card 300.